

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

SEQUOIA TECHNOLOGY, LLC,)	
)	
Plaintiff,)	
)	
v.)	Civil Action No. 18-1127-LPS-CJB
)	(CONSOLIDATED)
DELL INC., DELL TECHNOLOGIES)	
INC. (and its subsidiary EMC)	
CORPORATION (AKA DELL EMC)),)	
)	
Defendants.)	
)	
<hr/> RED HAT, INC.,)	
)	
Plaintiff,)	
)	
v.)	Civil Action No. 18-2027-LPS-CJB
)	
SEQUOIA TECHNOLOGY, LLC and)	
ELECTRONICS AND)	
TELECOMMUNICATIONS RESEARCH)	
INSTITUTE,)	
)	
Defendants.)	
)	
<hr/> SEQUOIA TECHNOLOGY, LLC,)	
)	
Counterclaim Plaintiff,)	
)	
v.)	
)	
RED HAT, INC. and INTERNATIONAL)	
BUSINESS MACHINES COPORATION,)	
)	
Counterclaim Defendants.)	
)	

REPORT AND RECOMMENDATION

In this consolidated action between, *inter alia*, Sequoia Technology, LLC (“Sequoia”) and Electronics and Telecommunications Research Institute (“ETRI”) and Red Hat, Inc. (“Red

Hat”), presently before the Court is the matter of claim construction. The Court recommends that the District Court adopt the constructions as set forth below.

I. BACKGROUND

A. Procedural Background

On July 31, 2018 and August 23, 2018, Sequoia filed Complaints in four different actions alleging infringement of United States Patent No. 6,718,436 (the “‘436 patent”); the Complaints were filed against four sets of Red Hat’s customers, which are, respectively, Defendants Dell, Inc., Dell Technologies, Inc. and EMC Corporation (in Civil Action No. 18-1127-LPS-CJB), Hewlett Packard Enterprise Co. (in Civil Action No. 18-1128-LPS-CJB), Hitachi Ltd. and Hitachi Vantara Corp. (in Civil Action No. 18-1129-LPS-CJB) and Super Micro Computer, Inc. (in Civil Action No. 18-1307-LPS-CJB).¹ On December 19, 2018, Red Hat filed a declaratory judgment action in Civil Action No. 18-2027-LPS-CJB against Sequoia and, thereafter, filed the operative First Amended Complaint (“FAC”) against both Sequoia and ETRI in that action. (Civil Action No. 18-2027-LPS-CJB, D.I. 16) Red Hat’s FAC seeks a declaratory judgment that it does not infringe the ‘436 patent and that the patent is invalid. (*Id.*)

These actions were thereafter all consolidated, with the lead case being Civil Action No. 18-1127-LPS-CJB. (D.I. 56) Chief Judge Leonard P. Stark has referred the consolidated cases to the Court to hear and resolve all pre-trial matters up to and including expert discovery matters. (See, e.g., D.I. 20)

¹ All citations herein, unless otherwise noted, are to the docket in the lead case, Civil Action No. 18-1127-LPS-CJB.

The parties filed their joint claim construction brief on June 11, 2020. (D.I. 153) On July 29, 2020, the Court conducted a *Markman* hearing by video conference. (D.I. 184 (hereinafter “Tr.”)).

B. Factual Background

Red Hat is a Delaware corporation and a “leading contributor to free and open source software[.]” (Civil Action No. 18-2027-LPS-CJB, D.I. 16 at ¶ 14) It manufactures Red Hat Enterprise Linux, or “RHEL,” for the commercial market. (*Id.* at ¶ 16) RHEL is accused of infringing the '436 patent. (*See, e.g., id.* at ¶ 27)

Sequoia is a Delaware limited liability company. (*Id.* at ¶ 7) ETRI is a South Korean research institution and the record owner of the '436 patent; it licenses the patent to Sequoia. (*Id.* at ¶¶ 8, 21)

The '436 patent is titled, “Method for Managing Logical Volume in Order to Support Dynamic Online Resizing and Software Raid and to Minimize Metadata and Computer Readable Medium Storing the Same[.]” ('436 patent, Title) The patent relates to “RAID,” or “Redundant Array of Independent Disks,” which is a “way of storing the same data to different locations of multiple hard disks [which] is usually utilized in a server with important data.” (*Id.*, col. 1:26-32) The invention described in the '436 patent relates to methods that work by constructing a “logical volume,” which is a “virtual disk drive,” out of “multiple physical disk drives[.]” (*Id.*, col. 1:24-26) The '436 patent uses a series of tables to keep track of where and how the data in a logical volume is located among the physical drives. (*Id.*, Abstract) By way of the disclosed methods, the patent aims to minimize the use of metadata and to “support dynamic online resizing” and RAID. (*Id.*) Further details regarding the '436 patent will be provided below in Section III.

II. STANDARD OF REVIEW

It is well-understood that “[a] claim in a patent provides the metes and bounds of the right which the patent confers on the patentee to exclude others from making, using, or selling the protected invention.” *Corning Glass Works v. Sumitomo Elec. U.S.A., Inc.*, 868 F.2d 1251, 1257 (Fed. Cir. 1989). Claim construction is generally a question of law, although subsidiary fact finding is sometimes necessary. *Teva Pharm. USA, Inc. v. Sandoz, Inc.*, 574 U.S. 318, 325-26 (2015).

The Court should typically assign claim terms their “‘ordinary and customary meaning[,]’” which is “the meaning that the term[s] would have to a person of ordinary skill in the art in question at the time of the invention, i.e., as of the effective filing date of the patent application.” *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312-13 (Fed. Cir. 2005) (citations omitted). However, when determining the ordinary meaning of claim terms, the Court should not extract and isolate those terms from the context of the patent; rather it should endeavor to reflect their “meaning to the ordinary artisan after reading the entire patent.” *Id.* at 1321; *see also Eon Corp. IP Holdings LLC v. Silver Spring Networks, Inc.*, 815 F.3d 1314, 1320 (Fed. Cir. 2016).

In proceeding with claim construction, the Court should look first and foremost to the language of the claims themselves, because “[i]t is a bedrock principle of patent law that the claims of a patent define the invention to which the patentee is entitled the right to exclude.” *Phillips*, 415 F.3d at 1312 (internal quotation marks and citations omitted). For example, the context in which a term is used in a claim may be “highly instructive.” *Id.* at 1314. In addition, “[o]ther claims of the patent in question, both asserted and unasserted, can . . . be valuable” in discerning the meaning of a particular claim term. *Id.* This is “[b]ecause claim terms are

normally used consistently throughout the patent, [and so] the usage of a term in one claim can often illuminate the meaning of the same term in other claims.” *Id.* Moreover, “[d]ifferences among claims can also be a useful guide[,]” as when “the presence of a dependent claim that adds a particular limitation gives rise to a presumption that the limitation in question is not present in the independent claim.” *Id.* at 1314-15.

In addition to the words of the claims, the Court should look to other intrinsic evidence. For example, the Court should analyze the patent specification, which “may reveal a special definition given to a claim term . . . that differs from the meaning [that term] would otherwise possess” or may reveal an intentional disclaimer of claim scope. *Id.* at 1316. Even if the specification does not contain such revelations, it “is always highly relevant to the claim construction analysis. Usually, it is dispositive; it is the single best guide to the meaning of a disputed term.” *Id.* at 1315 (internal quotation marks and citation omitted). That said, however, the specification “is not a substitute for, nor can it be used to rewrite, the chosen claim language.” *SuperGuide Corp. v. DirecTV Enters., Inc.*, 358 F.3d 870, 875 (Fed. Cir. 2004). And a court should also consider the patent’s prosecution history, if it is in evidence, because it “can often inform the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution[.]” *Phillips*, 415 F.3d at 1317.

Extrinsic evidence, “including expert and inventor testimony, dictionaries, and learned treatises[,]” can also “shed useful light on the relevant art[.]” *Id.* (internal quotation marks and citations omitted). Overall, while extrinsic evidence may be useful, it is “less significant than the intrinsic record in determining the legally operative meaning of claim language.” *Id.* (internal

quotation marks and citations omitted); *accord Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 981 (Fed. Cir. 1995).

In utilizing these resources during claim construction, courts should keep in mind that “[t]he construction that stays true to the claim language and most naturally aligns with the patent’s description of the invention will be, in the end, the correct construction.” *Renishaw PLC v. Marposs Societa’ per Azioni*, 158 F.3d 1243, 1250 (Fed. Cir. 1998).

III. DISCUSSION

The parties² presented eight disputed terms/term sets (“terms”) requiring construction. The Court will analyze these in turn below.

A. “disk partition”

The first disputed term, “disk partition” (sometimes appearing as “disk partitions”) appears in, *inter alia*, claims 1, 2, 3 and 8 of the '436 patent. Exemplary claim 1 recites:

1. A method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata, said method comprising steps of:
 - a) creating the logical volume by gathering *disk partitions* in response to a request for creating the logical volume in a physical storage space;
 - b) generating the metadata including information of the logical volume and the *disk partitions* forming the logical volume and storing the metadata to the *disk partitions* forming the logical volume;
 - c) dynamically resizing the logical volume in response to a request for resizing, and modifying the metadata on the *disk partitions* forming the logical volume; and
 - d) calculating and returning a physical address corresponding to a logical address of the logical volume by using mapping information of the metadata containing

² Unless otherwise noted, below when the Court refers to “Sequoia” it is referring to Sequoia and ETRI.

information of the physical address corresponding to the logical address;
 wherein the metadata includes,
 a disk partition table containing information of a *disk partition* in which the metadata is stored;
 a logical volume table for maintaining the information of the logical volume by storing duplicated information of the logical volume onto all *disk partitions* of the logical volume;
 an extent allocation table for indicating whether each extent in the *disk partition* is used or not used; and
 a mapping table for maintaining a mapping information for a physical address space corresponding to a logical address space which is a continuous address space equal in size of storage space to an entirety of said logical volume.

('436 patent, col. 12:16-48 (emphasis added)) The parties' proposed constructions are below:

Term	Sequoia's Proposal	Red Hat's Proposal
"disk partition"	"section of a disk"	"section of a disk that is a minimum unit of a logical volume"


(D.I. 153 at 7)

The parties' dispute is over whether (as Red Hat argues) a disk partition is the minimum unit of a logical volume. (*Id.* at 12, 16-17) Put differently, Sequoia argues that a logical volume can be constructed from something less than whole or entire disk partitions, while Red Hat argues that it can only be formed from whole or entire disk partitions. (*See id.* at 17; Tr. at 18-19) Before addressing the dispute, a bit of background about the patent-in-suit is in order.

The '436 patent, at a high level, is directed to, *inter alia*, "minimizing a size of metadata[.]" ('436 patent, Abstract & col. 4:55-61 (*cited in* Sequoia's *Markman* Presentation, Slide 7)) To achieve this aim, it utilizes a hierarchical process to store information. At the lowest level is the *extent*. An extent is "a minimum unit of space allocation to store

information[,]” and all extents are of uniform size. (‘436 patent, col. 7:1-3; *see also* Red Hat’s *Markman* Presentation, Slide 5; Tr. at 26) One level up from an extent is a *disk partition*, the term in dispute here. This is a “physical partition” of a disk. (‘436 patent, col. 6:58-59) Figure 2 illustrates how disk partitions are formed from physical disks:

FIG. 2



(*Id.*, FIG. 2) Here, “[physical d]isks 1, 2, 3, 4 are divided into four partitions and disks 5, 6, 7, 8 have one partition.” (*Id.*, col. 6:39-41) And the third and highest level is the *logical volume*. A logical volume is a “union of disk partitions and is extensible.” (*Id.*, col. 6:65) A logical volume “includes more than one disk partition” and its size is “resized in disk partition units.” (*Id.*, col. 6:60-63) An important distinction is that while disk partitions are physical sections of disks, a logical volume is “virtual” and extends across multiple disks. (See ‘436 patent, FIG. 2; Sequoia’s *Markman* Presentation, Slide 6) To keep track of where data is stored, the ‘436 patent, *inter alia*, correlates that data’s “physical address”—that is, its location in a physical disk—to its “logical

address”—that is, its “address” in relation to the logical volume it belongs to. ('436 patent, cols. 3:66-4:3)

With this background in mind, the Court turns to the parties’ present dispute. For three primary reasons, the Court concludes that Red Hat’s position here is the correct one.

First, the plain text of the claims favors Red Hat’s construction. Claim 1 (as do many other claims) describes how the method generates and modifies metadata of the “*disk partitions forming* the logical volume[.]” (*Id.*, col. 12:24-25, 28-29 (emphasis added)) It also describes “*creating* the logical volume by *gathering disk partitions[.]*” (*Id.*, col. 12:20 (emphasis added)) And claims 6 and 13 refer to “*disk partitions constructing* a logical volume[.]” (*Id.*, cols. 13:19-20, 15:7-8 (emphasis added)) These claim excerpts indicate that “disk partitions” are what are used to make up a logical volume—not smaller portions of disk partitions (i.e., a grouping of extents). (Tr. at 63-65) Put differently, the claims never indicate that logical volumes can be “*formed*” from anything less than a whole or entire disk partition, or that anything less than a whole or entire disk partition is “*gathered*” to create logical volumes. (*Id.* at 66; *see also* D.I. 153 at 11; Red Hat’s *Markman* Presentation, Slides 7-11)

Second, the specification also supports Red Hat’s construction. It recites that “[t]he present invention constructs a logical volume by using a *disk partition as a volume construction unit* so the present invention can minimize the size of metadata[.]” ('436 patent, cols. 11:66-12:1 (emphasis added)); *see GPNE Corp. v. Apple Inc.*, 830 F.3d 1365, 1371 (Fed. Cir. 2016) (“When a patent . . . describes the features of the ‘present invention’ as a whole, this description limits the scope of the invention[.]”) (internal quotation marks and citation omitted). And in another portion, where the patent describes the “three possible virtualizations of storage” that are “in accordance with a preferred embodiment of the . . . invention[,]” the specification also explains:

The first virtualization is a disk partition or physical partition. As described above, the disk partition is generated by an operating system's tool or utility. *The disk partition is a minimum unit of the logical volume.* A logical volume includes more than one disk partition. Therefore, a size of the logical volume is resized in disk partition units.

The second virtualization is the logical volume. The logical volume is a *union of disk partitions* and is extensible. It is named and provides continuous address space. The logical volume can be resized while the system is operating.

The third virtualization is an extent. The extent is continuous space having the same size. It is also a minimum unit of space allocation to store information. The size of an extent is fixed with each logical volume and it is decided at the creation of the logical volume. The size has to be an exponent of two.

After a disk partition is created using an operating system's tool or utility, a logical volume is *constructed with several disk partitions.*

('436 patent, cols. 6:55-7:9 (emphasis added)³; D.I. 178, ex. H at 5 (ETRI approvingly citing to this portion of the specification during an *inter partes* review ("IPR") proceeding regarding the '436 patent, for the proposition that "a '[disk] partition is a minimum unit of the logical volume'" and that "a 'logical volume is a union of disk partitions'"') (quoting '436 patent, col. 6:60-61, 6:65); *see also* '436 patent, cols. 5:64-67, 8:10-11, 8:42-43, 10:9-10) Just as in the claims, these descriptions of a disk partition as a "volume construction unit" or as "a minimum unit of the

³ To be sure, this portion of the specification is describing three virtualizations that are said to be in accord with a "preferred embodiment" of the invention. And because of this, Sequoia argues that the specification's explanation that "[t]he disk partition is a minimum unit of the logical volume[,]" should not be used to limit the scope of the claims. (D.I. 153 at 8 (citing '436 patent, col. 6:60-61); Tr. at 22, 35) But even though this statement is part of a discussion of a preferred embodiment, the particular language used by the patentee ("[t]he disk partition *is*") sure sounds a lot like a definition of a term that is meant to inform its meaning when used throughout the patent. And even if this is not meant to be a definitional statement *per se*, the bigger point is that the patent's description of a disk partition here is consistent with the patent's description of the term in the claims and in other portions of the specification—including such portions that do describe what the "present invention" is. (D.I. 153 at 12; Tr. at 70-71)

logical volume” strongly suggest that an entire or whole disk partition is used to form or construct a logical volume—not some lesser portion of a disk partition. (Tr. at 76)

Third, Red Hat’s position also appears to be consistent with statements that ETRI made during the IPR of the '436 patent. There, ETRI noted that “[w]hile the logical volume is formed from extents, extents are added or removed from a logical volume *at the level of disk partitions.*” (D.I. 178, ex. K at 5 (emphasis added)) And in the IPR, ETRI appeared to distinguish prior art (“Bridge” and “Williams”) in a manner that supports Red Hat’s position here. There it argued that “Bridge and Williams do not disclose or suggest gathering disk drives, the entity the petition alleges is a disk partition, to form a logical volume” and that “[i]nstead, extents in Bridge or physical partitions in Williams, both *subsets of disk drives*, . . . are gathered to form a logical volume.” (*Id.* at 2 (emphasis added)) Indeed, in denying institution of IPR, the United States Patent and Trademark Office’s (“USPTO”) Patent Trial and Appeal Board wrote that “Bridge’s logical volume is formed from *individual pieces* within one or more disk drives, *not from whole disk drives themselves.*” (*Id.*, ex. L at 21 (certain emphasis in original, certain emphasis added)) The Court need not determine whether ETRI’s statements in the IPR amounted to prosecution history disclaimer, (D.I. 153 at 14-15; Tr. at 30-31), in order to conclude that they are helpful to Red Hat’s position here (i.e., that logical volumes are not created by *subsets* of disk partitions, but instead from *whole* or *entire* disk partitions), (Tr. at 95).

The Court has also considered Sequoia’s arguments to the contrary, and does not find them persuasive.

For one thing, it is not entirely clear what Sequoia thinks *is* the minimum unit of a logical volume, if not a disk partition. At first during the *Markman* hearing, it seemed like Sequoia was suggesting that it is more appropriate to consider the *extent* the minimum unit of a logical

volume (since extents make up disk partitions, which in turn make up logical volumes). (Tr. at 20-21 ([Sequoia's Counsel:] "So look at the logical volume, Your Honor. If one were to ask what is that minimum unit of that logical volume, one of skill in the art or really any ordinary observer would say *the minimal unit of that logical volume is an extent*, because an extent makes up disk partitions, and that is what makes up a logical volume. And the patent says that, Your Honor. The patent says in certain embodiments, an extent is the minimum unit of a logical volume.")) (emphasis added))⁴ But at a later point in the hearing, Sequoia's counsel seemed to back away from this suggestion, stating that "Sequoia has never argued that extents create logical volumes." (*Id.* at 98) Instead, Sequoia's phraseology seemed to change, and it appeared to argue that a "portion of a [disk] partition" is the minimum unit of a logical volume—without clearly explaining what a "portion" of a partition *is*, if it is not a grouping of *extents* that amount to less than a whole or entire disk partition. (*Id.* at 102; *see also id.* at 41) In any event, the patent does not describe "portions of a disk partition" as making up logical volumes; instead, it makes clear that while a logical volume is formed from extents, extents are added or removed from a logical volume in units of whole or entire disk partitions.

Additionally, the Court does not agree with Sequoia that the patent's occasional references to a "whole partition of disks" or "whole disk partitions" is helpful to its case. ('436 patent, cols. 6:51, 11:23-30; Sequoia's *Markman* Presentation, Slide 16) Sequoia's argument

⁴ (*See also* D.I. 153 at 9 ("And the extents are the minimum units that make up the logical volumes."); Tr. at 21 ("[S]o the patent describes embodiments where the disk partition is a minimum unit of a logical volume, but it also describes embodiments where extents are the minimum units."); *id.* at 24 ("And so an extent really is the minimum unit [of a logical volume], not the disk partition."); *id.* at 25 ("[T]his idea of an extent being a minimum space of a unit of a corresponding [logical] volume was identified in our claim construction statement and the brief."))

here went as follows: (1) since the patent refers at two points to either “whole” partitions of disks or “whole” disk partitions; then (2) this implies the existence of something less than “whole” partitions (i.e., “partial” partitions); and (3) that, in turn, suggests that something less than a “whole” disk partition can form a logical volume. (Tr. at 18, 41-42, 99; *see also* D.I. 153 at 18) The Court declines to hop on to this logic train. The first of the two references at issue relates to Figure 2, *supra*, and describes how Volume 4 of the figure “includes the fourth partition of disks 1, 2, 3, 4 and the *whole partition of disks 5, 6, 7, 8[.]*” ('436 patent, col. 6:51-52 (emphasis added)) Thus, the use of “whole partition of disks” here simply seems like a way of distinguishing a disk partition that makes up an *entire* volume of a disk from a disk partition that is made up of only a *part* of the volume of a disk. (Red Hat’s *Markman* Presentation, Slide 26; D.I. 153 at 22; Tr. at 88-89) And the patent’s reference to “whole disk partitions” in column 11 is as follows:

The present invention can provide two methods: one method is to expand the RAID level to *whole disk partitions* including the new disk partition, and the other method is to expand the RAID level to *whole disk partitions* only on newly inserted data after the addition of new disk partition(s). In the latter method, data re-arrangement is not needed. Only newly inserted data are distributed considering the newly added disk partition(s).

('436 patent, col. 11:23-30 (emphasis added)) The Court sees nothing in this reference to “whole disk partitions” (made in the context of a discussion of expanding the RAID level) to suggest that logical volumes can be made up of something less than entire disk partitions.

For these reasons, the Court recommends that “disk partition” be construed to mean, “section of a disk that is a minimum unit of a logical volume.”

B. “logical volume”

The term “logical volume” appears in, *inter alia*, claims 1, 2 and 8 of the '436 patent.

The parties’ proposed constructions are as follows:

Term	Sequoia’s Proposal	Red Hat’s Proposal
“logical volume”	“a union of disk partitions that is extensible”	“extensible union of more than one disk partition, the size of which is resized in disk partition units”

(D.I. 153 at 23)

Similar to the previous term, the crux of the parties’ dispute is whether a logical volume can be formed from a sub-portion of a disk partition. Sequoia asserts that it can, while Red Hat asserts that it cannot. And the parties make similar arguments for this term (relying on similar evidence) as they did regarding “disk partition.” (*Id.* at 23-26 (Red Hat); *see also id.* at 25 (Red Hat incorporating its arguments as to “disk partition” for this term); Tr. at 104 (Sequoia), 106 (Red Hat))

Because the key, disputed portion of Red Hat’s construction (“the size of which is resized in disk partition units”) is drawn from the specification, ('436 patent, col. 6:62-63; *see also id.* at FIG. 13), and for the same reasons that the Court recommended adopting Red Hat’s construction as to “disk partition,” the Court also recommends adopting Red Hat’s construction as to “logical volume.”

Thus, the Court recommends that “logical volume” be construed to mean, “extensible union of more than one disk partition, the size of which is resized in disk partition units.”

C. “table”

The term “table” appears in, *inter alia*, claims 1, 2, 3 and 8 of the '436 patent. The parties’ proposed constructions are as follows:

Term	Sequoia's Proposal	Red Hat's Proposal
“table”	“structured list or map of data”	“data structure arranged in rows and columns”

(D.I. 153 at 29)

The parties appear to agree that, at a high level, a “table” contains data, and, that in some way, it can be used to relate one set of data to another set of data. (Tr. at 110, 128; Red Hat's *Markman* Presentation, Slide 76) The parties' dispute here is over whether a “table” necessarily needs to be arranged in “rows and columns” (with Red Hat arguing that it does, and Sequoia arguing that it does not). (D.I. 153 at 35; Tr. at 109, 113-14, 121, 128; Sequoia's *Markman* Presentation, Slide 36; Red Hat's *Markman* Presentation, Slide 94) And this is a challenging issue, since the patent does “not directly define what it means to store data . . . in the form of a ‘table[.]’” (D.I. 153 at 33)

The Court starts by focusing on Red Hat's proposed construction. For one thing, it does not seem that Red Hat's proposal can be word-for-word correct, because it facially requires a table to be arranged in “rows and columns” (plural). And yet Figure 3 of the patent, which all parties agree is a “table” (a “Metadata Table”), (Tr. at 116, 122; Red Hat's *Markman* Presentation, Slide 90), only depicts a *single* column (with five rows), (Tr. at 116; *see also* '436 patent, FIG. 3). (D.I. 153 at 30)⁵

Beyond that though, is Red Hat correct that the patent requires that “table” be at least arranged in a “row[-]and[-]column format[.]” (even if at times the table has only one row or one

⁵ In some contexts, perhaps, the use of a plural might encompass “a universe ranging from one to some higher number[.]” *Versa Corp. v. Ag-Bag Int'l Ltd.*, 392 F.3d 1325, 1330 (Fed. Cir. 2004); *see also* (D.I. 153 at 34). But here, Red Hat's use of the plural in its construction seems to detract from an accurate understanding of the term's meaning. (D.I. 153 at 38 n.24) And so the Court does not favor it for this reason.

column)? (D.I. 153 at 40; *see also* Tr. at 121) There is evidence in the record that both helps and harms Red Hat’s position in this regard.

On the one hand, it does appear that every embodiment of a “table” depicted in the patent could be described as employing a row-and-column structure. ('436 patent, FIGS. 3-7; *id.*, col. 5:27-44; Red Hat’s *Markman* Presentation, Slide 69)⁶ And Red Hat provided a declaration from its expert, Dr. Benjamin Goldberg, who opined that storing data in a table with rows and columns would be beneficial to the claimed invention, in that it would allow for “very fast access to data by being able to access immediately . . . a value at a particular row and column, rather than having to search through data as is generally required for a[] list or map structure.” (D.I. 154, ex. V at ¶ 15) Yet on the other hand, the patent never uses the words “row” or “column” to describe what a “table” is. (Tr. at 111-12) Moreover, although Red Hat cites to a number of technical dictionaries in support of its claim that a “table” necessarily has rows and columns,⁷ one of those dictionaries states only that a table “usually” has rows and columns—the implication being that some tables *do not have* them.⁸ And Sequoia, for its part, cites to its own

⁶ That said, some of the tables in the patent’s figures appear to depict rows and columns that are subsumed within another column or row. (D.I. 153 at 30; Tr. at 124; *see also* '436 patent, FIGS. 4-5)

⁷ See *A Glossary of Computing Terms* 15 (5th ed. 1987) (defining a table as “a data structure in the form of a rectangular arrangement of items in *rows and columns*”) (emphasis added) (D.I. 154, ex. S-1); *Computer Professional’s Dictionary* 319 (1990) (“[A] structure, consisting of a two-dimensional arrangement of *columns and rows*, in which data is stored.”) (emphasis added) (*id.*, ex. S-2); *Random House Webster’s Computer & Internet Dictionary* 544 (3d ed. 1999) (“[D]ata arranged in *rows and columns*. A *spreadsheet*, for example, is a table.”) (certain emphasis in original, certain emphasis added) (*id.*, ex. S-4); *Dictionary of Computer Science, Engineering, & Technology* 487 (2001) (“A collection of *rows* (or *tuples*) of data . . . with each *column* representing an attribute.”) (emphasis added) (*id.* ex. S-5).

⁸ See *The McGraw-Hill Illustrated Dictionary of Personal Computers* (4th ed. 1995) (“The items are *usually laid out in rows and columns* for reference or stored in memory as an array.”) (emphasis added) (D.I. 154, ex. S-3).

technical dictionaries, which include definitions of “table” that do not specifically include reference to rows and columns, such as “a data structure usually consisting of a list of entries, each entry being identified by a unique key and containing a set of related values[,]” *Microsoft Computer Dictionary* 510 (5th ed. 2002) (D.I. 154, ex. C), or a “collection of adjacent fields of data[,] [a]lso called an ‘array[,]’” ComputerLanguage, “table,” <https://www.computerlanguage.com/results.php?definition=table> (last visited Sept. 30, 2020) (D.I. 154, ex. D). (D.I. 153 at 39)

In the end, Red Hat is the party who is seeking the narrower, more limiting construction here. And with the evidence not decidedly coming down on its side of the ledger, the Court is unprepared to import Red Hat’s suggested limitation into the term’s construction. *See TI Grp. Automotive Sys. (N. Am.), Inc. v. VDO N. Am., L.L.C.*, 375 F.3d 1126, 1138 (Fed. Cir. 2004) (noting that the patentee is entitled to the “full breadth of claim scope supported by the words of the claims and the written description”).

As for Sequoia’s proposed construction, the Court does not find it helpful. For one thing, the Court does not see why it makes sense to refer to a “table” as a “list,” since: (1) the term “list” is never used in the patent specification, (D.I. 153 at 35; Tr. at 128); (2) the Court is not sure how use of a simple “list” would necessarily facilitate the relation of two sets of data to one another, (Red Hat’s *Markman* Presentation, Slide 83); and (3) during the *Markman* hearing, even Sequoia’s counsel seemed to back away from its reliance on this term, (Tr. at 117 (Sequoia’s Counsel: “Table 3 [i.e., Figure 3] says ‘map,’ which is our construction.”); *see also* D.I. 153 at 39 (Sequoia’s brief explaining how a “map” but not a “list” requires a relationship between data sets)) Moreover, with regard to Sequoia’s suggestion that a “table” can be a “map” of data, that appears to be an accurate statement, but one not particularly helpful to resolve the parties’ true

dispute here. The patent refers to how some logical volume managers in the prior art “ha[d] been using a fixed mapping method that uses a fixed convert function in converting the logical address . . . to a physical address of several underlying physical disk drives” but that other logical volume managers “don’t use a mapping function when a logical address is mapped to a physical address, but use a *table-based method for mapping* a logical address to physical address.” ('436 patent, col. 3:28-40 (emphasis added)) As Red Hat notes, (D.I. 153 at 36-37; Red Hat’s *Markman* Presentation, Slides 75-76), this excerpt suggests that a “table” can be a “map,”⁹ and that the process of “mapping” data could include the use of tables—but it also suggests that not all maps are tables and that “mapping” data does not necessarily *require* the use of tables. Because Sequoia’s proposal does not add much clarity to the meaning of “table,” the Court declines to adopt its construction.

Having resolved the disputes before it, but not being convinced that any of the proposed constructions are helpful at this stage, for now the Court recommends that “table” be afforded its plain and ordinary meaning.

D. “extent allocation table for indicating whether each extent in the disk partition is used or not used”

The next term, “extent allocation table for indicating whether each extent in the disk partition is used or not used,” appears in, *inter alia*, claims 1 and 8 of the '436 patent. The parties’ proposed constructions are as follows:

Term	Sequoia’s Proposal	Red Hat’s Proposal
------	--------------------	--------------------

⁹ The patent elsewhere refers to tables as a type of map. For example, Figure 3, which is a “metadata table[,]” includes various maps. ('436 patent, FIG. 3; *id.*, col. 5:28) Similarly, Figures 4-7 of the patent are described as a “[m]ap in the metadata table[.]” (*Id.*, col. 5:32-44; *see also* Tr. at 129) This gibus with the idea that a “map” is something that is used to correspond one set of data to another set of data, (Red Hat’s *Markman* Presentation, Slides 78-79), and that a “table” can be a somewhat more specific version of a “map.”

“extent allocation table for indicating whether each extent in the disk partition is used or not used”	No construction necessary.	“extent allocation table for indicating whether each extent in a disk partition is or is not storing information”
--	----------------------------	---

(D.I. 153 at 43)

The focus of the parties’ dispute is what it means for an extent to be “used or not used.”

In Sequoia’s view, an extent allocation table keeps track of whether or not an extent has been merely allocated to or assigned to a logical volume, and that this is what it means for the table to indicate whether the extent is “used or not used.” (D.I. 153 at 43-44, 47) But in Red Hat’s view, the extent allocation table must track whether each extent is actually *storing* information, and so “used or not used” means “is or is not storing information.” (D.I. 153 at 44-45; Red Hat’s *Markman* Presentation, Slides 111-12; Tr. at 141) For the three reasons set forth below, the Court agrees with Red Hat.

First, the claim language, read in context with certain of the specification’s text, provides some support for Red Hat’s position. The claim term at issue makes clear that the extent allocation table is meant to indicate whether each extent in a disk partition is “used or not used[,]” (‘436 patent, cols. 12:42-43, 13:56-57), and the specification says the same, (*id.*, col. 7:65-67 (“[t]he extent allocation map **74** of FIG. 3 uses one bit per each extent in the disk partition and it represents *usage* of a corresponding extent”) (emphasis added)). And we know from elsewhere in the specification that an extent is “a minimum unit of space allocation *to store information[,]*” (*id.*, col. 7:2-3 (emphasis added))—i.e., that an extent’s basic purpose, when it comes to the claimed inventions, is for storing information. (See D.I. 153 at 44, 49) So it follows that when the claim requires the extent allocation table to show whether an extent is

“used” or not, that “usage” designation should be one that sheds light on whether the extent is actually storing information.

Second, a disclosure in Chang-Soo Kim, et al., *Volume Management in SAN Environment*; Proc. of the Eighth Int’l Conf. on Parallel and Distributed Sys., 500-505 (2001) (“*Volume Management*”) (D.I. 178, ex. G), a paper authored by the named inventors and cited in the patent (such that it is part of the intrinsic record), ('436 patent at 1), also bolsters Red Hat’s argument. *Volume Management* describes an embodiment of the '436 patent:

The SANtopia [system, an early embodiment of the patent’s claims] allocates and deallocates the storage space in extent unit. The extent is a contiguous storage space that consists of multiple blocks. An extent may be *used* for both *normal data* and *metadata*. *The SANtopia gives two bits to the allocation bitmap for an extent in order to distinguish these usages of an extent.* The value 00 is given to an extent for the free space, 01 is for an inode, 10 is for a directory entry and 11 is for a data extent.

Volume Management at 502 (emphasis added). This description is consistent with Red Hat’s proposed construction, in that it indicates that an extent is “used” when it is storing information—either storing normal data (as indicated by the value 11) or metadata (as indicated by the values 01 or 10). (D.I. 153 at 46) Free space (which contains no information) is denoted by a value of “00.” (*Id.*; *see also* D.I. 154, ex. V at ¶ 28)

Third, Sequoia’s proposal would render superfluous the extent allocation map referenced in a preferred embodiment. As part of this embodiment, “[t]he extent allocation map **74** of FIG. 3 uses one bit per each extent *in the disk partition* and it represents usage of a corresponding extent.” ('436 patent, col. 7:65-67 (emphasis added)) “[T]he disk partition” implies a single disk partition that these extents are “in”—that is, that they have been allocated to or assigned to. Yet under Sequoia’s construction, this table would (nonsensically) record the same value for each extent—because this table would be showing *allocation* and not *storage*. (*Id.*) This just raises the question: Why have

this table at all, if it would record the same value throughout? There does not appear to be a good answer to this—another point in favor of Red Hat’s construction.

For these reasons, the Court recommends that “extent allocation table for indicating whether each extent in the disk partition is used or not used” be construed to mean “extent allocation table for indicating whether each extent in a disk partition is or is not storing information.”

E. “mirror data table for validating data”

The next disputed term, “mirror data table for validating data,” appears only in dependent claim 3, which recites:

3. The method as recited in claim 1, wherein the metadata further includes a *mirror data table for validating data* located on several disk partitions corresponding to each mirror by maintaining the mirror data table in disk partitions constructing each mirror and storing the mirror data table to each corresponding mirror in an overlapped manner.

(*Id.*, col. 12:53-58 (emphasis added)) The parties’ proposed constructions are as follows:

Term	Sequoia’s Proposal	Red Hat’s Proposal
“mirror data table for validating data”	“table, containing bits, for verifying data that is duplicated”	“table containing bits indicating the validity of data duplicated in multiple partitions”

(D.I. 153 at 50)

The parties have two key disputes as to this term: (1) whether the term should be construed to mean that the bits within the table “indicat[e] the validity of [the] data duplicated” (as Red Hat proposes), or whether (as Sequoia advocates) the table itself (instead of its constituent bits) may perform this task, (*id.* at 51, 53); and (2) whether the table itself (as Red Hat argues) indicates a “conclusion” about “the validity and the status of the mirrors[,]” (Tr. at 147), or whether (as Sequoia urges) the table need only include *data from which one could*

derive such a conclusion about whether or not the data is valid, (*id.*). The Court finds that Red Hat has the better argument in both regards.

The parties' disputes turn in large part on the following portion of the specification, which both sides look to in divining this term's meaning, (D.I. 153 at 50, 52; Sequoia *Markman* Presentation at Slides 50-53; Red Hat's *Markman* Presentation, Slides 132-33, 137), and which describes the only disclosed embodiment of the claimed mirror data table:

The mirror write consistency map **73** is maintained only when the logical volume is constructed by mirror type and it is duplicated and stored at each mirror. For example, if disk partition **1** and disk partition **2** are mirrored, then MWC map **73** of the disk partition **1** is duplicated as MWC map **73** of the disk partition **2**. It is used for maintaining precise information when a malfunction occurs in one of the mirrored disk partitions.

...

The MWC map **73** of FIG. **3** is constructed with three bits per one extent when mirror data is two (there are two mirrors)[.] Therefore, the first three bits are about the first extent. *The first bit is validity of original data, the second bit is validity of the first mirror data, and the third bit is validity of the second mirror data.*

As an example, the first mirror of disk partition **1** is disk partition **2**, and the second mirror is disk partition **3**. The first *bit* of MWC map **73** is *validity of the first extent* of the disk partition **1**, the second *bit* is *validity of the first extent* of the disk partition **2**, and the third *bit* is *validity of the first extent* of the disk partition **3**. The fourth bit is validity of the second extent of the disk partition **1**.

The MWC map **73** is duplicated and stored at disk partitions **1**, **2** and **3**. It is for *verifying validity* even if any one or two of the disk partitions **1**, **2**, **3** malfunction(s).

('436 patent, col. 7:30-37, 49-64 (emphasis added))

As to the first dispute, to be sure (as Sequoia notes), the claim term itself states that the "mirror data table" is what does the work of validating data. (D.I. 153 at 51) But the

specification excerpt above describes *how* the table does so: on a *bit-by-bit* basis, not by using the table as a whole in some way.

As to the second dispute, the specification’s text again supports Red Hat: it describes how “[t]he first bit *is* validity of original data, the second bit *is* validity of the first mirror data,” and so on. (‘436 patent, col. 7:52-53 (emphasis added)) That is, this description demonstrates that the bits tell us whether the data *is* or *is not* valid, without the need for any further steps to be taken to derive that answer.¹⁰

Therefore, the Court recommends that “mirror data table for validating data” be construed as “table containing bits indicating the validity of data duplicated in multiple partitions.”

F. “a method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata”

The next term, “a method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata,” appears in the preambles of, *inter alia*, claims 1 and 8. (‘436 patent, cols. 12:16-18, 13:31-33) The parties dispute is about whether the preambles are limiting (with Red Hat asserting that they are, and Sequoia asserting that they are not). The Court agrees with Red Hat that they are.

¹⁰ In their briefing, the parties at one point had a running dispute about whether Red Hat’s proposed construction required the table at issue to include the use of an “algorithm, code or function.” (D.I. 153 at 53) Sequoia argued that the “[claimed] table, in and of itself does not include . . . the algorithm, code or function for validating.” (*Id.*) Yet in response, Red Hat clarified that its proposed construction was not meant to require that the claimed table includes such an algorithm, code or function. (*Id.* at 55 n.30) Thus, there is no dispute as to this issue.

Additionally, for a time it looked like there was a dispute about the “in multiple partitions” portion of Red Hat’s construction. (D.I. 153 at 51, 53) But Red Hat ultimately confirmed that, like Sequoia, it intended that any construction of the term should require only that data need be duplicated in the original partition and at least one other, not that data must be duplicated in a total of three or more partitions. (*Id.* at 55 n.30)

The question of whether language in a preamble constitutes a claim limitation is a question of law. *Rotatable Techs. LLC v. Motorola Mobility LLC*, 567 F. App'x 941, 943 (Fed. Cir. 2014). On the one hand, “[w]hile it is true that preamble language is often treated as nonlimiting in nature, it is not unusual for [the United States Court of Appeals for the Federal Circuit] to treat preamble language as limiting[.]” *Bicon, Inc. v. Straumann Co.*, 441 F.3d 945, 952 (Fed. Cir. 2006). Generally, “a preamble limits the invention if it recites essential structure or steps, or if it is necessary to give life, meaning, and vitality to the claim.” *Catalina Mktg. Int'l, Inc. v. Coolsavings.com, Inc.*, 289 F.3d 801, 808 (Fed. Cir. 2002) (internal quotation marks and citation omitted). Relatedly, a preamble may also be construed as limiting when: (1) the claim limitations in the body of the claim “rely upon and derive antecedent basis from the preamble[.]” *Eaton Corp. v. Rockwell Int'l Corp.*, 323 F.3d 1332, 1339 (Fed. Cir. 2003); or (2) the patent applicants clearly relied on the benefits or features recited therein as patentably significant, *Catalina*, 289 F.3d at 808-09. On the other hand, when the claim body recites a “structurally complete invention” and the preamble language is used merely to state the “purpose or intended use of the invention[.]” the preamble is generally not treated as limiting the scope of the claim. *Id.* at 808.¹¹

The first indication that the preambles are limiting is that they contain antecedent basis for the claim terms “logical volume” and “metadata,” which are found in the body of the claims.

¹¹ There is no “litmus test” for determining whether preamble language is limiting. *Catalina*, 289 F.3d at 808. Rather, whether such language is limiting is assessed in regard to “the facts of each case in light of the claim as a whole and the invention described in the patent.” *Storage Tech. Corp. v. Cisco Sys., Inc.*, 329 F.3d 823, 831 (Fed. Cir. 2003); *see also Catalina*, 289 F.3d at 808 (“Whether to treat a preamble as a limitation is a determination resolved only on review of the entire . . . patent to gain an understanding of what the inventors actually invented and intended to encompass by the claim.”) (internal quotation marks, brackets and citation omitted).

(D.I. 153 at 57; *see also id.* at 58; Red Hat’s *Markman* Presentation, Slides 152-53) This is, at the very least, a strong indication that Red Hat’s position is correct. *See Pacing Techs., LLC v. Garmin Int’l, Inc.*, 778 F.3d 1021, 1024 (Fed. Cir. 2015) (“Because the preamble terms ‘user’ and ‘repetitive motion pacing system’ provide antecedent basis for and are necessary to understand positive limitations in the body of claims in the ‘843 patent, we hold that the preamble to claim 25 is limiting.”); *In re Mobile Telecomms. Techs., LLC*, 265 F. Supp. 3d 454, 475-76 (D. Del. 2017).

But more significantly in Red Hat’s favor is that: (1) the preamble phrase states that the method/medium at issue is for “support[ing] dynamic online resizing and minimizing a size of metadata[,]”; and (2) the ‘436 patent repeatedly states that “the present invention” is all about doing both of those things. Preamble language “will limit the claim if it recites not merely a context in which the invention may be used, but *the essence of the invention* without which performance of the recited steps is nothing but an academic exercise.” *Boehringer Ingelheim Ventmedica, Inc. v. Schering-Plough Corp.*, 320 F.3d 1339, 1345 (Fed. Cir. 2003) (emphasis added). And here, the specification repeatedly hammers home the point that “managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata” is the “essence of the invention.” (See, e.g., ‘436 patent, cols. 1:10-21 (“The present invention relates to a method for managing a logical volume for minimizing a size of metadata and dynamic resizing, and a computer-readable recording medium . . . embodying the method[.]”), 3:66-4:3 (“It is, therefore, an object of the present invention to provide a logical volume manager . . . using a minimum space for metadata and supporting online dynamic resizing.”); *see also id.* at

Abstract & cols. 4:55-58, 5:4-12, 11:66-12:5)¹² Given how important these aims are to the invention of the '436 patent, this too dictates that the preambles (which reference the aims) are limiting. *See Boehringer Ingelheim*, 320 F.3d at 1345 (concluding that, in a claim where the preamble recited that a “method for growing and isolating” a certain virus was claimed, that “growing” and “isolating” were relevant claim limitations, because they “[were] not merely circumstances in which the method may be useful, but instead are the *raison d'être* of the claimed method itself”) (emphasis in original); *see also Cadence Pharms., Inc. v. Paddock Labs. Inc.*, 886 F. Supp. 2d 445, 460-61 (D. Del. 2012). And because the concept of “minimizing” metadata is not repeated in the body of the claims, finding the preambles to be limiting is important here, as it ensures that the scope of the claims lines up neatly with what the inventors repeatedly emphasized was an object of the invention.¹³ (Red Hat’s *Markman* Presentation, Slide 158)

Thus, the Court recommends that the preamble phrase “a method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata” be construed as limiting.

¹² Indeed, the specification distinguishes the present invention over prior art systems that required large amounts of metadata. ('436 patent, col. 3:41-63)

¹³ Sequoia argues that the preambles cannot be limiting because “‘resizing’ . . . is repeated in the body of the claim language, whereas ‘minimizing [a size of metadata]’ is not.” (D.I. 153 at 59) But as Red Hat points out, “[i]f the preamble were required to be repeated in the [claim] body to be limiting, then a preamble would *never* be a meaningful limitation.” (Red Hat’s *Markman* Presentation, Slide 159 (emphasis in original)) And, relatedly, in plenty of cases the Federal Circuit has found preamble language to be limiting where it was not repeated in the claim’s body. *See, e.g., Deere & Co. v. Bush Hog, LLC*, 703 F.3d 1349, 1357-58 (Fed. Cir. 2012); *Boehringer Ingelheim*, 320 F.3d at 1344-45.

G. “computer-readable recording medium”

The next disputed term, “computer-readable recording medium” or “CRM” appears in the preamble of claim 8, which recites:

8. A computer-readable recording medium storing instructions for executing a method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata, said method comprising the steps of: . . .

('436 patent, col. 13:30-33 (emphasis added)) The parties' constructions are:

Term	Sequoia's Proposal	Red Hat's Proposal
“computer-readable recording medium” (“CRM”)	No construction needed.	“transitory or non-transitory computer-readable recording medium”

(D.I. 153 at 60)

The parties' dispute is over whether this term can include “transitory” mediums.¹⁴ Red Hat says it can; Sequoia says it cannot. The Court agrees with Red Hat, which has the better of the arguments here, in light of the available record.

The specification discusses a “computer readable medium” in the following way:

The above-mentioned method of the present invention can be stored in a computer readable medium *including* compact disc read only memory (CDROM), random access memory (RAM), floppy disk, hard disk, and magneto-optical disk.

¹⁴ The parties spent portions of their respective briefing focusing on whether this claim is patent-ineligible under 35 U.S.C. § 101 and in light of the Federal Circuit's decision in *In re Nuijten*, 500 F.3d 1346, 1351-52 (Fed. Cir. 2007), on the ground that it embraces non-statutory subject matter. (D.I. 153 at 61, 65-66, 68; Red Hat's *Markman* Presentation, Slides 165-66; Sequoia's *Markman* Presentation, Slide 63) But at this phase of the case, the task before the Court is to construe the claim term at issue, not to determine eligibility. So the parties' arguments as to eligibility are better raised at a case dispositive stage. *See Acceleration Bay LLC v. Activision Blizzard, Inc.*, 324 F. Supp. 3d 470, 486-87 (D. Del. 2018) (deciding the eligibility of a claim to a “computer readable medium” at the summary judgment stage, where the term at issue had previously been construed during claim construction).

(436 patent, col. 11:36-39 (emphasis added)) Sequoia accurately notes that this list of examples of a CRM “does not include any transitory signals.” (D.I. 153 at 66; *see also* Sequoia’s *Markman* Presentation, Slide 62) But the use of the word “including” in the above excerpt indicates that other examples of a CRM can and do exist. *See Yodlee, Inc. v. Plaid Techs., Inc.*, C.A. No. 14-1445-LPS, 2016 WL 204372, at *13 (D. Del. Jan. 15, 2016). So this excerpt does not state that claim 8’s CRM, which stores the requisite instructions for executing the method, cannot encompass transitory mediums (i.e., signals or waves). Nor does any other portion of the specification limit a CRM to non-transitory mediums.

Moreover, the record indicates that a person of ordinary skill in the art (“POSITA”) would, at the time of the priority date of the '436 patent (i.e., December 2001), have understood “computer-readable recording medium” to embrace transitory forms. For support on this point, Red Hat points again to Dr. Goldberg’s declaration. There, Dr. Goldberg explains that he has reviewed excerpts from 34 different prior art patents and patent applications dating from 2000 and 2001 (which Red Hat has also provided to the Court); these references clearly indicate that a claim directed to a “computer-readable recording medium” in the relevant time frame would ordinarily have been understood to encompass transitory mediums, such as carrier waves. (D.I. 154, ex. V at ¶ 35; *id.*, ex. U)¹⁵

¹⁵ Dr. Goldberg’s conclusion is consistent with that of other Courts who have examined this issue. *See Acceleration Bay LLC*, 2017 WL 6508715, at *3-4 (construing the term “computer readable medium” found in patents with priority dates in July 2000 as embracing transitory media, including carrier waves, even though the specification never expressly mentioned carrier waves, because the parties agreed that a carrier wave is one type of transitory “computer readable medium” and in light of a USPTO definition of “computer readable medium” that stated that the term typically covers transitory signals); *cf. Ex Parte Mewherter*, Appeal 2012-007692, 2013 WL 3291360, at *2-5 & nn. 2, 5 (PTAB May 8, 2013) (construing the term “machine readable storage medium” (a term that is equivalent to a “computer-readable medium”), albeit using a “broadest reasonable interpretation” standard, and concluding that it embraced transitory signals, when “[the] Specification fail[ed] to limit expressly the term . . . to

For its part, Sequoia also submitted a declaration in support of its position from its expert, Zaydoon Jawadi. But Mr. Jawadi simply opines, in fairly brisk and conclusory fashion, that a “POSITA would not understand that the CRM claim applied herein encompasses any transitory signals.” (D.I. 154, ex. B at ¶ 37) He does not attempt to grapple with the prior art cited by Dr. Goldberg, nor does he provide much more in the way of detail in support of his conclusion. So the Court does not give Mr. Jawadi’s opinion great weight.

For the above reasons, the Court sees no reason to adopt Sequoia’s narrower construction. It recommends that “computer-readable recording medium” be construed as “transitory or non-transitory computer-readable recording medium.”

H. Claim 8

The next disputed issue also relates to claim 8, which in full recites:

8. A computer-readable recording medium storing instructions for executing a method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata, said method comprising the steps of:

- a) creating the logical volume by gathering disk partitions in response to a request for creating the logical volume in a physical storage space;
- b) generating the metadata including information of the logical volume and the disk partitions forming the logical volume and storing it the metadata to the disk partitions forming the logical volume;
- c) dynamically resizing the logical volume in response to a request for resizing, and modifying the metadata on the disk partitions forming the logical volume; and
- d) calculating and returning a physical address corresponding to a logical address of the logical volume by using mapping

exclude signals, carrier waves, etc.” and citing to various pieces of extrinsic evidence from 2001-03 that indicated that the term was broad enough to include both non-transitory and transitory media) (internal quotation marks omitted)).

information of the metadata containing information of the physical address corresponding to the logical address;

wherein the metadata includes,

a disk partition table containing information of a disk partition in which the metadata is stored;

a logical volume table for maintaining the information of the logical volume by storing duplicated information of the logical volume onto all disk partitions of the logical volume;

an extent allocation table for indicating whether each extent in the disk partition is used or not used; and

a mapping table for maintaining a mapping information for a physical address space corresponding to a logical address space which is a continuous address space equal in size of storage space to an entirety of said logical volume.

('436 patent, col. 13:30-62 (emphasis added))

Red Hat argues that this claim is invalid as indefinite for improperly mixing apparatus and method classes of statutory subject matter. (D.I. 153 at 73; Red Hat's *Markman* Presentation, Slides 181-93) For the reasons set forth below, the Court concludes that Red Hat has not met its burden to demonstrate that the claim is invalid. *See Bayer Pharma AG v. Watson Labs., Inc.*, Civil Action No. 12-1726-LPS-CJB, 2014 WL 4954617, at *3 (D. Del. Sept. 30, 2014).

In *IPXL Holdings, L.L.C. v. Amazon.com, Inc.*, 430 F.3d 1377 (Fed. Cir. 2005), the Federal Circuit invalidated a claim directed to “the *system of claim 2* [including an input means] wherein the predicted transaction information comprises both a transaction type and transaction parameters associated with that transaction type, and *the user uses the input means* to either change the predicted transaction information or accept the displayed transaction type and transaction parameters.” 430 F.3d at 1384 (citation omitted). The *IPXL Holdings* Court concluded that the claim “recite[d] . . . both [a] system . . . and a method for using that system[.]”

Id. Due to the hybrid nature of the claim, the Court explained that “it is unclear whether infringement . . . occurs when one creates a system that allows the user to change the predicted transaction information or accept the displayed transaction, or whether infringement occurs when the user actually uses the input means to change transaction information or uses the input means to accept a displayed transaction.” *Id.* (emphasis added)

The rule articulated in *IPXL Holdings* has repeatedly been recognized to be a narrow one. *See Bayer Pharma AG*, 2014 WL 4954617, at *6 (citing cases). To that end, courts have explained that the rule does not apply to claims containing language simply describing a system as well as the capabilities of the claimed system; rather, the rule applies to claims describing a system that also require the user of the recited system to take specific action. *Compare In re Katz Interactive Call Processing Patent Litig.*, 639 F.3d 1303, 1318 (Fed. Cir. 2011) (applying *IPXL Holdings* to invalidate system claims as indefinite where “the language used in [the claims at issue] . . . is directed to user actions, not system capabilities”), *with Microprocessor Enhancement Corp. v. Texas Instruments, Inc.*, 520 F.3d 1367, 1375 (Fed. Cir. 2008) (finding that claim was not indefinite under *IPXL Holdings* where it was “clearly limited to a pipelined processor possessing the recited structure and capable of performing the recited functions”). Accordingly, courts should “focus on whether the claim language is directed to user actions rather than system capabilities.” *H-W Tech., LC v. Overstock.com. Inc.*, 973 F. Supp. 2d 689, 696 (N.D. Tex. 2013); *see also Bayer Pharma. AG*, 2014 WL 4954617, at *6.

Here, claim 8 recites a “computer-readable recording medium” having instructions for executing a particular method, where that method includes the steps of “creating the logical volume” and “generating the metadata[.]” ('436 patent, col. 13:30-37) Nothing in the claim language actually requires the user to do anything. Instead, the claim simply requires that the

medium must contain all instructions that would permit the medium to perform the steps further set out in the claim. (D.I. 153 at 78-79)

To conclude, the Court does not recommend that claim 8 should be found invalid for being a mixed method/apparatus claim.

IV. CONCLUSION

For the foregoing reasons, the Court recommends that the District Court adopt the following constructions/recommendations:

1. “disk partition” should be construed to mean “section of a disk that is a minimum unit of a logical volume”;
2. “logical volume” should be construed to mean “extensible union of more than one disk partition, the size of which is resized in disk partition units”;
3. “table” should be afforded its plain and ordinary meaning;
4. “extent allocation table for indicating whether each extent in the disk partition is used or not used” should be construed to mean “extent allocation table for indicating whether each extent in a disk partition is or is not storing information”;
5. “mirror data table for validating data” should be construed to mean “table containing bits indicating the validity of data duplicated in multiple partitions”;
6. “a method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata” should be construed as limiting;
7. “computer-readable recording medium” should be construed to mean “transitory or non-transitory computer-readable recording medium”; and
8. Claim 8 should not be found indefinite for being a mixed method/apparatus claim.

This Report and Recommendation is filed pursuant to 28 U.S.C. § 636(b)(1)(B), Fed. R. Civ. P. 72(b)(1), and D. Del. LR 72.1. The parties may serve and file specific written objections within fourteen (14) days after being served with a copy of this Report and Recommendation. Fed. R. Civ. P. 72(b)(2). The failure of a party to object to legal conclusions may result in the

loss of the right to *de novo* review in the district court. *See Sincavage v. Barnhart*, 171 F. App'x 924, 925 n.1 (3d Cir. 2006); *Henderson v. Carlson*, 812 F.2d 874, 878-79 (3d Cir. 1987).

The parties are directed to the Court's Standing Order for Objections Filed Under Fed. R. Civ. P. 72, dated October 9, 2013, a copy of which is available on the District Court's website, located at <http://www.ded.uscourts.gov>.

Dated: October 1, 2020

Christopher J. Burke
Christopher J. Burke
UNITED STATES MAGISTRATE JUDGE